

# An improved current potential method for fast computation of stellarator coil shapes

Matt Landreman<sup>\*</sup>

*Institute for Research in Electronics and Applied Physics,  
University of Maryland, College Park, MD, 20742, USA*

(Dated: February 10, 2017)

Several fast methods for computing stellarator coil shapes are compared, including the classical NESCOIL procedure [Merkel, Nucl. Fusion 27, 867 (1987)], its generalization using truncated singular value decomposition, and a Tikhonov regularization approach we call REGCOIL in which the squared current density is included in the objective function. Considering W7-X and NCSX geometries, and for any desired level of regularization, we find the REGCOIL approach simultaneously achieves lower surface-averaged and maximum values of both current density (on the coil winding surface) and normal magnetic field (on the desired plasma surface). This approach therefore can simultaneously improve the free-boundary reconstruction of the target plasma shape while substantially increasing the minimum distances between coils, preventing collisions between coils while improving access for ports and maintenance. The REGCOIL method also allows finer control over the level of regularization, it preserves convexity to ensure the local optimum found is the global optimum, and it eliminates two pathologies of NESCOIL: the resulting coil shapes become independent of the arbitrary choice of angles used to parameterize the coil surface, and the resulting coil shapes converge rather than diverge as Fourier resolution is increased. We therefore contend that REGCOIL should be used instead of NESCOIL for applications in which a fast and robust method for coil calculation is needed, such as when targeting coil complexity in fixed-boundary plasma optimization, or for scoping new stellarator geometries.

---

<sup>\*</sup> [mattland@umd.edu](mailto:mattland@umd.edu)

## I. INTRODUCTION

Stellarators offer the promise of highly stable, disruption-free, steady-state fusion plasmas with no power recirculated for current drive, but there are also significant challenges for the stellarator concept related to the complicated magnetic field coils. The magnetic field coils are among the most expensive components in any magnetic fusion system, and the non-planar coils of a stellarator are particularly difficult to design and fabricate due to the need for precise three-dimensional shaping. The ability of non-planar coils to shape the magnetic field is constrained by the requirement that the coils cannot collide (overlap). Even without collisions, the minimum coil-to-coil separation in stellarators is typically small, limiting access between the coils for diagnostics and heating systems, and for blanket maintenance in a reactor. Any advances in stellarator coil design can potentially have a significant impact on the cost and feasibility of fusion energy.

For recent optimized stellarators such as W7-X [1, 2] and HSX [3], the coil design has been part of a two-stage optimization process. In the first stage, the plasma boundary shape is varied to optimize physics properties such as neoclassical transport and magnetohydrodynamic (MHD) stability, with a fixed-boundary solution of the MHD equilibrium equations at each iteration, and no specific coil shapes computed. In the second stage, coil shapes are calculated that approximately produce the ‘target’ plasma shape resulting from stage 1. This two-stage approach is advantageous both because it is computationally efficient (there is no need to compute free-boundary equilibria from a new coilset at each iteration in stage 1), and because good flux surface quality is more likely than in a single-stage variation of coil shapes.

For this multi-stage optimization procedure, a central tool of stellarator coil design has been the NESCOIL code [4]. In the NESCOIL approach, one first defines a coil winding surface on which the coils will lie, which might be done by expanding the target plasma boundary outward by some uniform or nonuniform distance. One then considers a ‘current potential’  $\Phi$ , which is a scalar stream function for a divergence-free surface current  $\mathbf{K}$  flowing on this coil surface:

$$\mathbf{K} = \mathbf{n} \times \nabla \Phi, \quad (1)$$

where  $\mathbf{n}$  is the unit surface normal. NESCOIL computes the current potential that best produces the target plasma shape, in a least-squares sense. Remarkably, this problem can be solved by solving a single linear system of modest size (typically  $\leq 144 \times 144$ ), making the procedure fast and robust. (The typical time for a well-resolved calculation, which is dominated by matrix assembly, is on the order of a few seconds on a modern CPU.) Once the continuous current potential is obtained, the shapes of discrete coils are defined by taking a finite number of contours of  $\Phi$ . There is essentially only one parameter requiring routine adjustment, the number of Fourier modes retained in  $\Phi$ , making the code minimally complicated to use. On the other hand, NESCOIL does not permit the direct inclusion of other important optimization criteria or constraints, such as the maximum magnetic field at the coils, minimum coil radius of curvature, or minimum separation between coils. For this reason, nonlinear optimization tools for stellarator coil design have been developed, including ONSET[5], COILOPT[6, 7] (used to design NCSX [8]), and COILOPT++[9].

While it is certain that such nonlinear optimization with engineering constraints will be essential to the design of future stellarator facilities, the NESCOIL approach remains widely used, for several reasons. First, due to NESCOIL’s speed, robustness, and relatively minimal complexity, it is often used for initial studies of new stellarator configurations, before a detailed engineering design is conducted. Recent examples of such use include [10–14]. Second, the convergence of a nonlinear coil optimizer is expedited by using a good initial guess for the coil shapes, and NESCOIL has been used to provide this initial condition [12]. Third, NESCOIL is sometimes called in the first stage of plasma optimization, to guide the fixed-boundary plasma optimization towards configurations consistent with realistic coils [15]. Since the maximum feasible coil-to-plasma distance (or equivalently the coil complexity at fixed coil-to-plasma distance) is a strong function of the plasma shape [16], this application is highly important. NESCOIL’s speed and reliability make it well suited for use inside the phase-1 optimization iterations in this manner. For all these applications, NESCOIL is uniquely robust in part because the optimization problem it solves is convex, so NESCOIL can never be ‘stuck’ in a local minimum that is not the global minimum, unlike most other coil optimization methods.

In this paper, we compare NESCOIL to two variants of the original method. The first variant is a truncated singular value decomposition (TSVD) option that was added to the NESCOIL code during the NCSX design process [15, 17]. In the second variant, which we call REGCOIL (regularized NESCOIL) for brevity, the average squared current density is added as an optimization criterion, in a manner that preserves the linearity, convexity, and speed of the original NESCOIL method. We will demonstrate that for both W7-X and NCSX geometries, the REGCOIL approach systematically yields better coil shapes than the other two methods. This is shown first in the sense that the maximum and average-squared current density is lower for a given level of error in the magnetic field on the plasma boundary. (Or equivalently, for any given maximum or average-squared current density, the error in the magnetic field is lower.) We also demonstrate these advantages using free-boundary reconstructions with discrete coils, showing coilsets from REGCOIL can more accurately reconstruct the target plasma shape in practice than coilsets from the other methods,

while simultaneously having greater coil-to-coil separation. Furthermore, we point out a shortcoming of the original NESCOIL and TSVD methods, that the resulting coil shapes can be sensitive to the arbitrary choice of how the coil surface is parameterized, meaning the code user must pay attention to this parameterization, and potentially causing problems with use of NESCOIL inside an optimizer. In contrast, the REGCOIL formulation yields results that are independent of the choice of parameterization. Thus, we will show that REGCOIL has multiple advantages compared to the original NESCOIL and TSVD variants, at comparable computational cost.

In most coil optimization work to date, departures from the target magnetic field are weighted uniformly in the objective function (in a sense that will be made precise in section III A.) However it would be preferable to place greater weight on the magnetic field errors that couple strongly to important physics properties, such as magnetic field errors with a helicity that resonates with the rotational transform somewhere in the plasma, driving magnetic islands. Methods have been proposed for incorporating this plasma sensitivity information into coil design [18, 19]. Here, for simplicity, we take the usual approach of uniform weighting of errors, to focus on the comparison of regularization methods.

## II. ILL-POSED PROBLEMS AND REGULARIZATION

Before describing the three algorithms in detail, one should note an important fact: the problem of finding currents that produce a given plasma shape is fundamentally ill-posed. Since the magnetic field from a coil decays with distance, and since oppositely directed currents can give nearly cancelling contributions to the Biot-Savart integral, certain aspects of the coil shapes are effectively undetermined. Given one set of coils that produce a certain plasma shape, one can in principle always add an out-and-back segment to any coil without significantly affecting the plasma shape, as in figure 1. Thus, there cannot be a unique solution for the coil shapes, unless extra criteria are provided to disallow configurations like that of figure 1.b. Analogous ill-posed problems arise in many other fields [20], often in the context of least-squares problems, another example from fusion being tomography [21].




FIG. 1. (Color online) Finding coil shapes that produce a given magnetic field is an ill-posed problem since  $O(1)$  changes to the coil shape, such as the change from (a) to (b) here, can produce negligible change to the magnetic field in the region of interest. Here, the difference between (a) and (b) of the magnetic field in the plasma region becomes arbitrarily small as the distance  $\delta \rightarrow 0$  due to near cancellation in the Biot-Savart law.

In any numerical scheme that attempts to solve a fundamentally ill-posed problem, typically a linear system will arise for which the matrix has a large condition number [20], with a solution that does not converge as numerical resolution is increased. Ill-posed problems can be ‘regularized’ in a variety of ways by imposing some additional constraints, such that the underlying problem becomes well-posed, and the solution is no longer underdetermined. One regularization approach can be to look for the solution only within a particular subspace, such as the space of polynomials of degree  $p$  for some finite  $p$ . Another regularization approach can be to impose a constraint on the solution norm. The different algorithms considered below for computation of stellarator coil shapes essentially differ in their choice of regularization technique.

In any regularization method for an ill-posed problem, one is always left with a degree of arbitrariness in the solution, represented by some parameter that reflects how much regularization to impose. As the original problem has an infinite number of solutions, this arbitrariness in the solution is unavoidable. The solution to any regularized ill-posed problem typically transitions from smooth to noisy (or vice-versa) as this regularization parameter is varied. This behavior is familiar to anyone who has fit data with a polynomial, in which case the regularization parameter is the polynomial degree  $p$ : raising  $p$  reduces differences between the data and fit, but if  $p$  is too large there is overfitting. There is no universal rule to choose the regularization parameter; the parameter often must be chosen by trial-and-error, or by using additional information or constraints outside the problem formulation. Indeed, we will find that the three coil-computation algorithms each has one regularization parameter that the user is forced to choose.

The existence of an arbitrary regularization parameter also reflects the fact that coil design is an optimization problem with multiple competing objectives. The coils should produce a plasma with the desired shape (or more to

the point, a plasma with good physics properties), but the coils should also be far from each other, with large radius of curvature, etc. A multi-objective optimization problem [22] generally does not have a unique solution, in the sense that an improvement can typically be achieved for one objective if a sacrifice is made in another objective. A unique solution only exists if some additional information is supplied to weigh the competing objectives. The regularization parameter is precisely this additional information, a relative weight given to the coil shape complexity compared to the plasma shape reconstruction.

### III. LINEAR CURRENT POTENTIAL METHODS FOR COMPUTING STELLARATOR COIL SHAPES

#### A. NESCOIL

In the NESCOIL method, an objective function  $\chi_B^2$  is defined in terms of  $B_{\text{normal}}$ , the normal component of the magnetic field on the desired (target) plasma surface, by integrating its square over the plasma surface:

$$\chi_B^2 = \int d^2a B_{\text{normal}}^2. \quad (2)$$

The fixed-boundary target plasma configuration has  $B_{\text{normal}} = 0$  exactly, but any attempt to produce this desired plasma shape using external coils will result in some finite error  $B_{\text{normal}} \neq 0$ , leading to  $\chi_B^2 > 0$ .

Next, note that the current potential  $\Phi$  is generally not single-valued, but by integrating (1) around the coil surface poloidally or toroidally it follows that

$$\Phi(\theta', \zeta') = \Phi_{sv}(\theta', \zeta') + \frac{G\zeta'}{2\pi} + \frac{I\theta'}{2\pi}, \quad (3)$$

where  $\Phi_{sv}$  is single-valued,  $\theta'$  and  $\zeta'$  are any poloidal and toroidal angles on the coil surface, and  $G$  and  $I$  are the currents linking the coil surface poloidally and toroidally. Here and throughout, we use primes to indicate coordinates on the coil surface, with unprimed coordinates indicating the plasma surface. The secular terms in (3) involving  $G$  and  $I$  are considered known, since  $G$  and  $I$  are determined by the desired coil topology and target plasma configuration. (For example, saddle coils are defined by  $G = I = 0$  so the coil filaments do not link the coil surface. For modular coils,  $I = 0$ , and  $G$  can be determined from the target plasma equilibrium by applying Ampere's Law to a toroidal loop at the plasma edge.) The goal becomes determination of the unknown single-valued term  $\Phi_{sv}$ .

The linearity of Ampere's Law and (1) mean that  $B_{\text{normal}}$  can be written as a superposition of the contributions from various currents:

$$B_{\text{normal}} = B_{\text{normal}}^{\text{plasma}} + B_{\text{normal}}^{\text{external}} + B_{\text{normal}}^{GI} + B_{\text{normal}}^{\text{sv}}\{\Phi_{sv}\}, \quad (4)$$

where  $B_{\text{normal}}^{\text{plasma}}$  is the component arising due to currents in the plasma,  $B_{\text{normal}}^{GI}$  is the component arising due to the secular  $G$  and  $I$  terms in the current potential (3),  $B_{\text{normal}}^{\text{sv}}$  is the component arising due to the single-valued term  $\Phi_{sv}$  in (3), and  $B_{\text{normal}}^{\text{external}}$  is the component arising due to any other external coils that might exist, such as planar toroidal field coils. The operator  $B_{\text{normal}}^{\text{sv}}$  can be derived from the Biot-Savart law, so it is linear, and an explicit formula is given in appendix A. In essence, one wants to determine the  $\Phi_{sv}$  that approximately drives an equal and opposite  $B_{\text{normal}}$  to the other terms in  $B_{\text{normal}}$ ,  $B_{\text{normal}}^{\text{plasma}} + B_{\text{normal}}^{\text{external}} + B_{\text{normal}}^{GI}$ .

The optimization problem is discretized by writing  $\Phi_{sv}$  as a finite Fourier series with amplitudes  $\Phi_j$ :

$$\Phi_{sv}(\theta', \zeta') = \sum_j \Phi_j \begin{pmatrix} \sin \\ \cos \end{pmatrix}_j (m_j \theta' - n_j \zeta'), \quad (5)$$

where  $j$  indexes both the poloidal and toroidal mode numbers  $(m, n)$  as well as indicating whether a sine or cosine phase is used. For stellarator-symmetric geometries, only the sine phase is needed.

Determination of the  $\Phi_j$  that minimizes  $\chi_B^2$  then has the form of a linear least-squares problem. Detailed expressions for the method are given in appendix A. As with any linear least-squares problem, the solution can be obtained using the normal equations,  $QR$  decomposition, or singular value decomposition. The number of unknowns ( $\Phi_j$ ) retained is typically 24-144, i.e. quite modest. Thus, the procedure is robust and fast: the current potential is computed by solving a single linear system, with no iteration required. The solution is unique, so there is no danger of getting trapped in a local minimum that is not the global minimum, as could be the case for non-convex optimization problems.

In the NESCOIL approach, the ill-posed problem is regularized through the truncation of the Fourier expansion (5), i.e. only optimizing the current potential within a restricted space of functions. Typically modes  $m_j$  and  $n_j$  are included which satisfy  $|m_j| \leq N_F$  and  $|n_j| \leq N_F$  for some number  $N_F$  in the range 3-8.



## B. Truncated SVD

Any linear least-squares problem such as the one solved by NESCOIL may be solved using a singular value decomposition (SVD). The SVD of a  $m \times n$  matrix  $A$  is  $A = U\Sigma V^T$  where  $U$  is a  $m \times m$  orthogonal matrix,  $V$  is a  $n \times n$  orthogonal matrix, and  $\Sigma$  is a  $m \times n$  diagonal matrix, with the diagonal entries  $\sigma_j$  non-negative and in decreasing order. The columns of  $U$  and  $V$  are known as the left and right singular vectors respectively, and the  $\sigma_j$  are the singular values. For the linear least-squares problem of determining the  $n$ -element vector  $x$  that minimizes  $|Ax - b|^2$  for a  $m$ -element vector  $b$ , it can be shown that the solution is  $x = V\Sigma^+U^Tb$ . Here  $\Sigma^+$  is the  $n \times m$  diagonal matrix with  $j$ th diagonal element equal to  $1/\sigma_j$  if  $\sigma_j > 0$ , or to 0 if  $\sigma_j = 0$ . The matrix  $V\Sigma^+U^T$  used to compute the solution is called the pseudoinverse of  $A$ . For NESCOIL,  $x$  corresponds to  $\Phi_j$ , and  $A$  and  $b$  are given in (A12).

A standard approach to regularizing ill-posed least-squares problems is using a truncated SVD (TSVD) [20]. In this method, all diagonal elements of  $\Sigma^+$  are replaced with 0 after row  $\ell$ , where  $\ell$  becomes the regularization parameter. That is, only  $\ell$  singular values and singular vectors are retained in the pseudoinverse. The idea behind this method is that the decreasing sequence of singular values (with corresponding singular vectors) corresponds to patterns in the solution that contribute less and less to the residual. Beyond a certain point in this sequence, these patterns in the solution are effectively undetermined, so their amplitude is set to 0. A TSVD solution to the minimization problem (A12) was added to the NESCOIL code during the NCSX design process [15, 17].

## C. REGCOIL

Another widely-used technique for regularizing ill-posed least-squares problems in other fields is Tikhonov regularization, sometimes also called ridge regression [20, 23]. In Tikhonov regularization, the original least-squares problem  $\min_x \{|Ax - b|^2\}$  is augmented with a second quadratic term:

$$\min_x \{|Ax - b|^2 + \lambda |L(x - x_0)|^2\}, \quad (6)$$

for some scalar  $\lambda$ , matrix  $L$ , and vector  $x_0$ . The idea is to regularize the problem by finding a solution  $x$  for which the difference from  $x_0$  has a small norm, in a sense given by  $L$ . In the common case that  $x_0 = 0$  and  $L$  is the identity matrix, it is just the usual 2-norm of the solution that is minimized. The regularization parameter is  $\lambda$ , representing how much importance to place on minimizing the solution norm compared to the competing goal of minimizing the original residual.

The REGCOIL method for computing stellarator coil shapes is the application of Tikhonov regularization to the least-squares problem from NESCOIL, taking the regularization term  $|L(x - x_0)|^2$  to be a physically meaningful quantity: the surface-average-squared current density:

$$\chi_K^2 = \int d^2a' K(\theta', \zeta')^2, \quad (7)$$

where  $K = |\mathbf{K}|$  reflects the inverse distance between coils. An explicit expression for  $\chi_K^2$  in terms of the unknown  $\Phi_j$  is given in appendix A, where it can be seen that  $\chi_K^2$  indeed has the form of the regularization term in (6). In other words, REGCOIL finds the current potential that minimizes the objective function  $\chi^2$  defined by

$$\chi^2 = \chi_B^2 + \lambda \chi_K^2, \quad (8)$$

with regularization parameter  $\lambda$ . Crucially, since  $K^2$  is a quadratic function of  $\Phi_{sv}$ , the problem of finding the  $\Phi_{sv}$  that minimizes (8) remains a *linear* least-squares problem, with a unique solution that can be obtained rapidly using normal equations,  $QR$  decomposition, or a SVD. Other constraints that one might include in a detailed engineering optimization, such as the maximum  $K$ , radius of curvature of the coils, or maximum magnetic field strength  $B$  on the coils, would require an iterative optimization, and would potentially introduce the possibility of multiple local minima in the objective function.

Although the REGCOIL method had been implemented in a U.S. version of the NESCOIL code previously, we are unaware of any publications in which the method has been described or demonstrated. Results shown below were generated using an independent parallelized implementation of REGCOIL, using OpenMP parallelization for assembling the matrices, (typically the most expensive step in any of these three methods), and using multi-threaded LAPACK routines for solving the least-squares problem.

#### IV. VARIATION WITH LEVEL OF REGULARIZATION

We will apply the three coil calculation algorithms to two geometries, NCSX [8] and W7-X [1, 2], shown in figure 2. In both cases the plasma surface is taken from a fixed-boundary VMEC equilibrium [24] free of coil ripple. For NCSX this equilibrium is the one known as LI383. For each geometry, the coil surfaces used are surfaces on which the actual experimental coils lie.

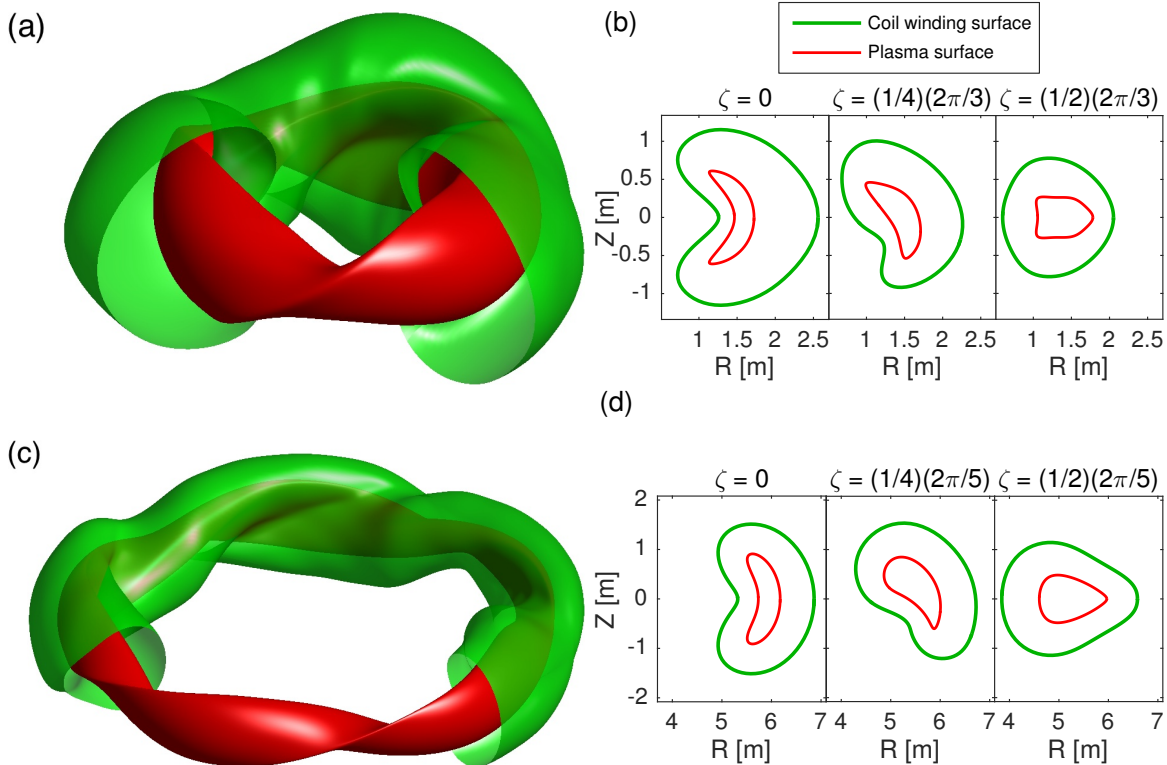


FIG. 2. (Color online) Shapes of the plasma and coil winding surfaces used in this work. (a)-(b) NCSX (c)-(d) W7-X.

As a first application, we illustrate how the three methods behave as the level of regularization is varied in each approach. While the specific mechanism of regularization is different in the three methods, several trends are similar: as the level of regularization is increased, the coil shapes become less complicated, the current density  $K$  decreases, and the residual  $|B_{\text{normal}}|$  increases. Figures 3, 4, and 5 demonstrate these trends for the NCSX geometry for the three algorithms. For NESCOIL, the level of regularization is increased by lowering  $N_F$ , the maximum poloidal and toroidal Fourier mode number. For the TSVD method, the level of regularization is increased by including fewer singular values (with corresponding singular vectors) in the pseudoinverse. For REGCOIL, the level of regularization is increased by increasing  $\lambda$ . The qualitative trends in the three methods are similar. In particular,  $K$  decreases as  $|B_{\text{normal}}|$  increases. Correlated with decreasing  $K$  is decreasing complexity of the coil shapes given by the contours of the current potential. Ideally one wants both low  $K$  and low  $|B_{\text{normal}}|$ , but there is a tradeoff between these competing objectives for all three algorithms.

While a finite Fourier series is used for the current potential in all three methods, increasing the Fourier resolution  $N_F$  causes the current potential from NESCOIL to diverge, whereas results for the TSVD and REGCOIL methods converge. The divergence of NESCOIL can be seen in figure 3, while the convergence of REGCOIL is demonstrated in figure 6. In this figure, the tradeoff curve of  $\chi_B^2$  versus  $\chi_K^2$  is plotted for the full range of  $\lambda$  from 0 to  $\infty$ , varying  $N_F$ . More and more of the curve can be resolved as  $N_F$  is increased. The  $\lambda = 0$  end of the curve corresponds to NESCOIL, so it diverges to  $\chi_K^2 \rightarrow \infty$  as  $N_F \rightarrow \infty$ . For any finite positive  $\lambda$ , if the Fourier resolution  $N_F$  is low enough, more regularization is provided by the Fourier truncation than by the physical  $\chi_K^2$  regularization, leading to higher  $\chi_B^2$  than the converged values; such behavior can be seen in figure 6 for the  $N_F = 8, 12$ , and 16 curves. In practice, the interesting solutions lie on the vertical part of the L-shaped curve, for  $\chi_B^2 > 10^{-4} \text{ T}^2 \text{ m}^2$ . Thus, we find  $N_F \sim 12$  is typically plenty for well converged results. Convergence behavior of the TSVD method is analogous to REGCOIL. For results shown throughout this paper, for the surface integrations detailed in appendix A, grid resolutions of  $128 \times 128$  points on the plasma and coil surfaces were used for all three methods, and it was verified that differences were

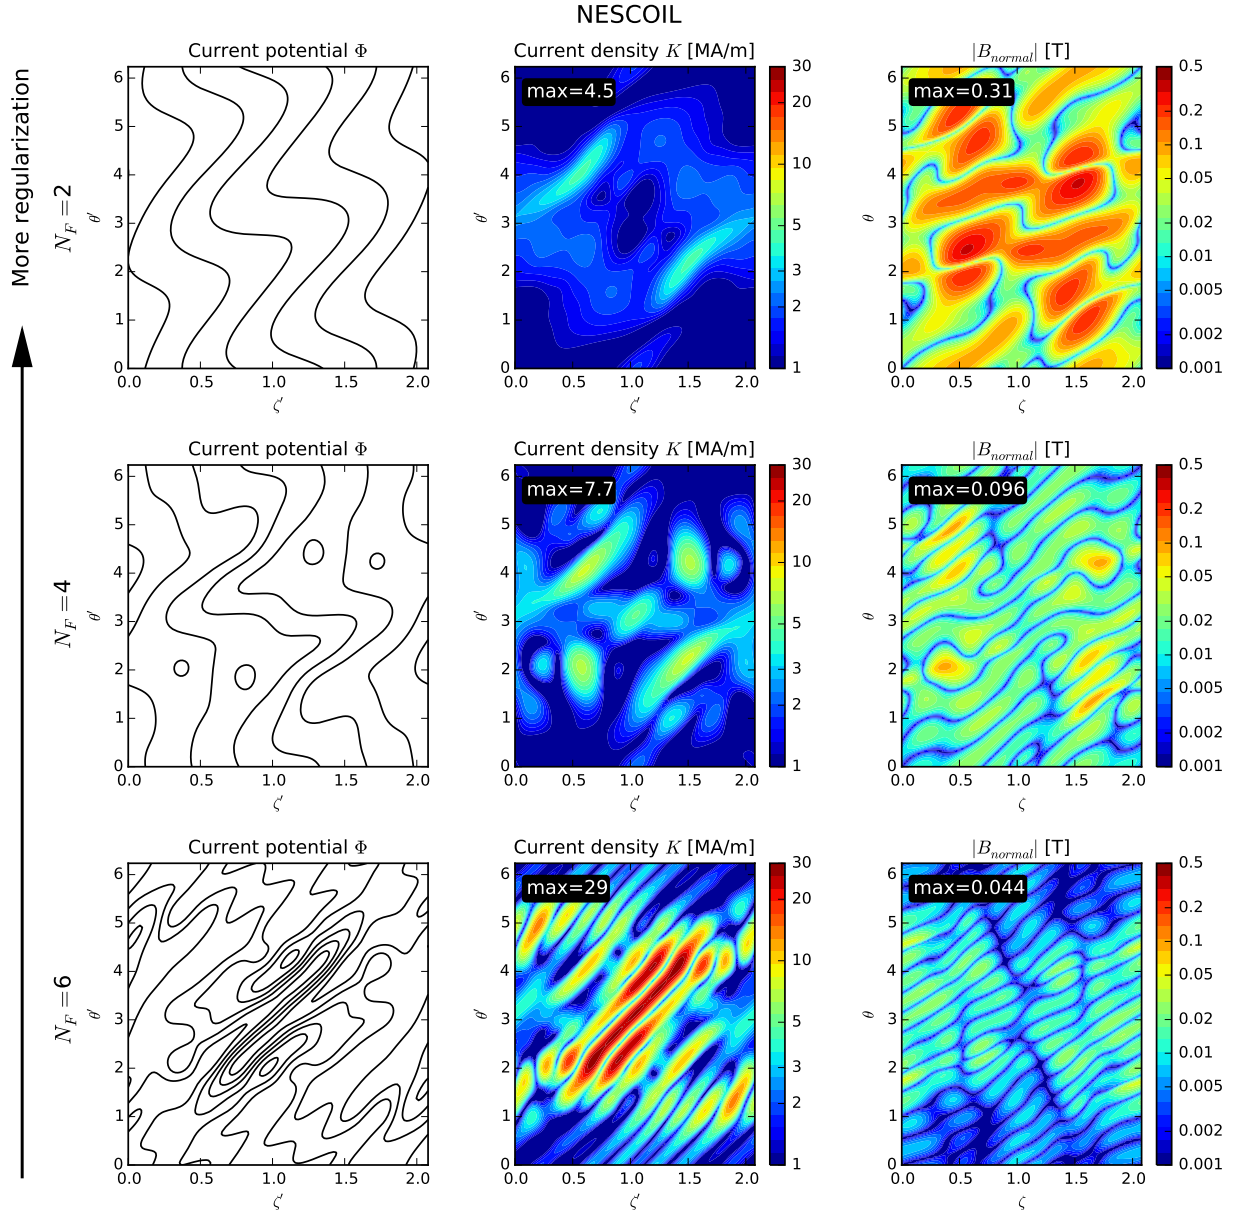


FIG. 3. (Color online) Application of NESCOIL to the NCSX geometry, showing trends with varying regularization parameter  $N_F$ . Regularization is provided in NESCOIL by truncating the Fourier expansion of the current potential to modes with  $|m| \leq N_F$  and  $|n| \leq N_F$ , so larger  $N_F$  corresponds to weaker regularization. As the level of regularization is increased, the coil shapes given by contours of the current potential become less complicated, and the current density  $K$  on the coil surface decreases, but the residual  $|B_{\text{normal}}|$  on the target plasma surface increases. Poloidal and toroidal angles on the coil surface are  $(\theta', \zeta')$ , and corresponding angles on the plasma surface are  $(\theta, \zeta)$ .

negligible if a resolution of  $256 \times 256$  was used. Computing the  $N_F = 12$  curve of figure 6, representing 100 unique coilsets with different values of  $\lambda$ , required 7 seconds on 1 node of the NERSC Cori computer.

## V. COMPARISON OF THE METHODS

Comparing figures 3-5, it can be seen that each row of the REGCOIL results has comparable or better  $K$  and  $B_{\text{normal}}$  to the corresponding rows for the other two methods. This behavior is systematic, as shown in figure 7. Parts (a)-(b) of the figure show the averaged  $B_{\text{normal}}$  error  $\chi_B^2$  and squared current density  $\chi_K^2$  for the three algorithms and two

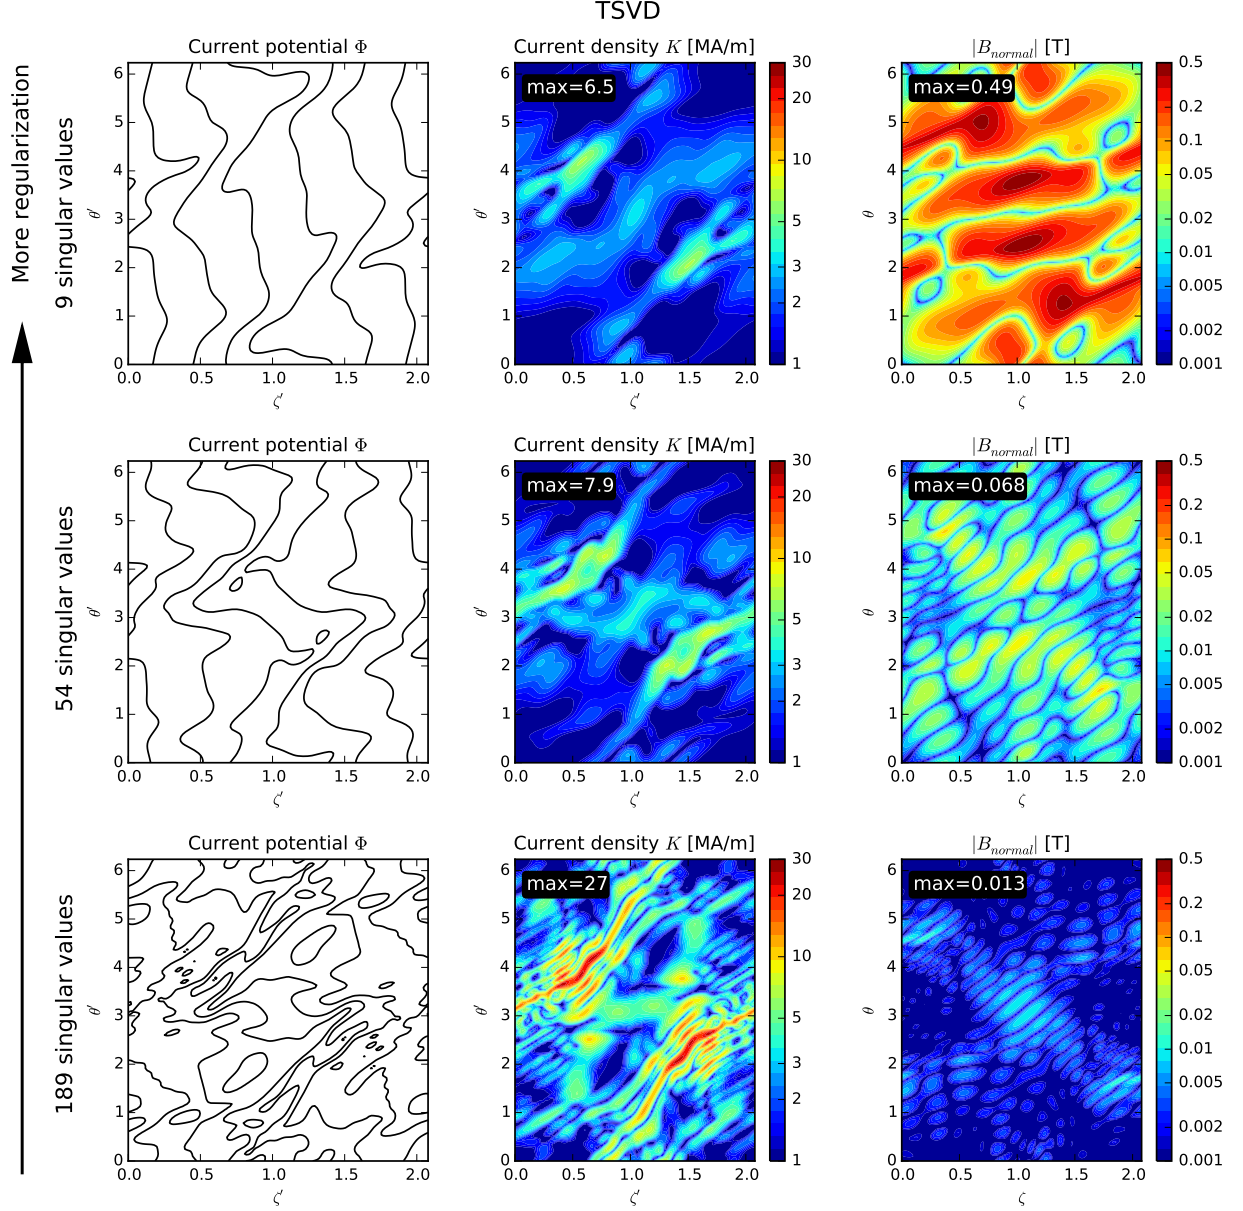


FIG. 4. (Color online) Application of the truncated singular value decomposition (TSVD) method to the NCSX geometry, showing trends with varying regularization. Regularization is provided in the TSVD approach by truncating the number of singular values retained in the pseudoinverse, so more singular values corresponds to less regularization. As the level of regularization is increased, the coil shapes given by contours of the current potential become less complicated, and the current density  $K$  on the coil surface decreases, but the residual  $|B_{normal}|$  on the target plasma surface increases.

geometries, scanning the amount of regularization for each algorithm over a wide range. For NESCOIL, results are displayed for  $N_F = 1, 2, 3, \dots$  (Using different limits for the number of poloidal and toroidal Fourier modes results in worse performance, i.e. higher  $\chi_K^2$  for given  $\chi_B^2$  and vice-versa, corresponding to points above and to the right of those plotted.) For the TSVD method, the level of regularization is scanned by varying the number of singular values and vectors retained in the pseudoinverse from  $1, 2, 3, \dots$ . For REGCOIL,  $\lambda$  is scanned from 0 to  $\infty$ . When  $\lambda \rightarrow \infty$ , corresponding to minimizing  $\chi_K^2$  with no contribution from  $\chi_B^2$ , a finite and regular solution results for the current potential, corresponding to current filaments separated by the maximum possible average distance. When  $\lambda \rightarrow 0$ , corresponding to minimizing  $\chi_B^2$  with no contribution from  $\chi_K^2$ , the current potential diverges as numerical resolution is increased, just as for NESCOIL. The solutions displayed in figures 3-5 are indicated with bold symbols. Unsurprisingly, REGCOIL consistently obtains the best (lowest)  $\chi_B^2$  for given  $\chi_K^2$ , and lowest  $\chi_K^2$  for given  $\chi_B^2$ , since



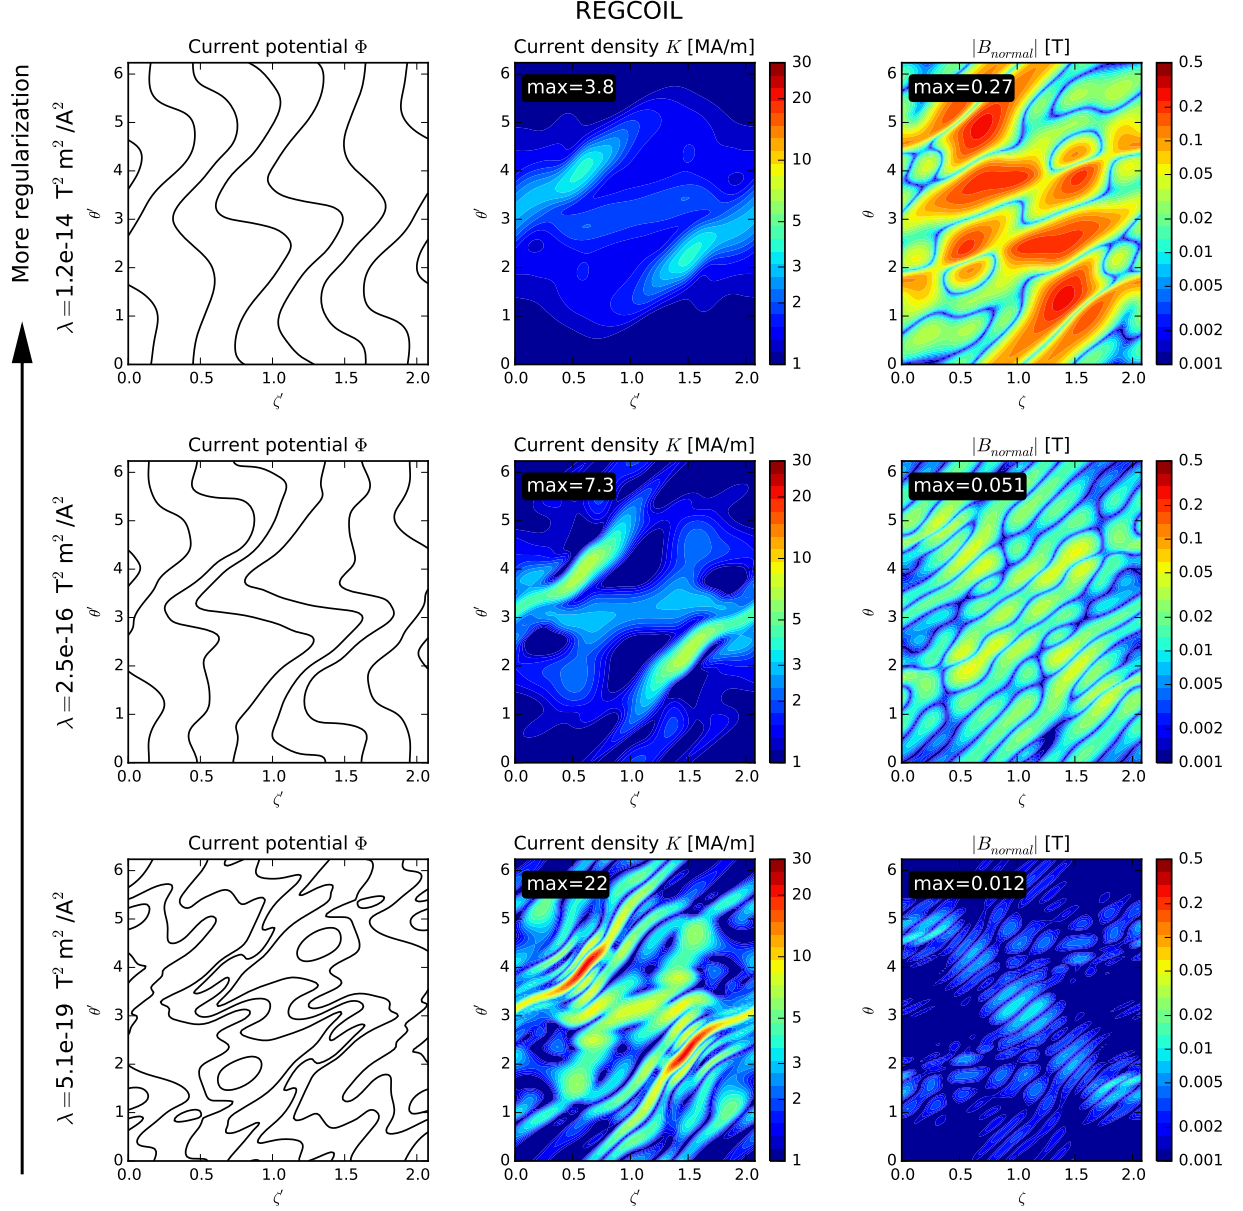


FIG. 5. (Color online) Application of REGCOIL to the NCSX geometry, showing trends with varying regularization parameter  $\lambda$ . As the level of regularization is increased, the coil shapes given by contours of the current potential become less complicated, and the current density  $K$  on the coil surface decreases, but the residual  $|B_{\text{normal}}|$  on the target plasma surface increases.

the REGCOIL method was constructed precisely to compute this optimum. In the terminology of multi-objective optimization, the NESCOIL and TSVD solutions are ‘dominated’ by those of REGCOIL, corresponding to the former being above and to the right of the latter. Indeed, since  $\chi_B^2$  and  $\chi_K^2$  are convex functions of  $\Phi_j$ , minimization of the weighted sum (8) directly yields the so-called ‘Pareto frontier’ of all possible non-dominated solutions to the multi-objective optimization problem [22], so no other algorithm could possibly obtain lower  $\chi_B^2$  and  $\chi_K^2$  simultaneously. The TSVD method tends to perform a bit worse than NESCOIL for high amounts of regularization, but performs better than NESCOIL for low levels of regularization.

Next, figures 7.c-d show the tradeoff in maximum rather than surface-average-squared  $|B_{\text{normal}}|$  and  $K$ . The REGCOIL method is not directly minimizing these criteria, there is no fundamental reason why the REGCOIL method must always outperform NESCOIL and the TSVD methods by these measures. Nonetheless, the figures show that the NESCOIL and TSVD results lie above and to the right of the REGCOIL results, aside from a few NCSX TSVD solutions at




FIG. 6. (Color online) Convergence of the REGCOIL method, using W7-X geometry. As the Fourier resolution  $N_F$  is increased, more and more of the tradeoff curve is resolved, corresponding to solutions of lower regularity.

very low regularization on the far right. (The current potential at this low level of regularization is highly contorted so these solutions are uninteresting.) Evidently, the fact that REGCOIL minimizes the surface-average-squared quantities means it tends to do a good job minimizing the maximum quantities, even though it is not a direct optimization for the latter.

One way to understand the improved performance of REGCOIL over NESCOIL is the following. In NESCOIL, one only optimizes the current potential within an artificially circumscribed space, the space of functions that have a very short Fourier expansion. By contrast, in the REGCOIL approach, the current potential can be optimized over a larger space of functions, with longer Fourier expansions. Although some current potentials with long Fourier expansions correspond to impractical coils, some current potentials with moderately long Fourier expansions do correspond to practical coils. By removing the constraint on the Fourier series length, both  $\chi_B^2$  and  $\chi_K^2$  can be improved simultaneously. Since the Fourier series length is a less physically meaningful criterion than  $\chi_K^2$  (which corresponds to the inverse distance between coils), it is advantageous to sacrifice the former for the latter.

Another noteworthy observation is that for a given residual  $B_{\text{normal}}$ , the TSVD method tends to produce the most contorted coil shapes. This behavior is particularly visible in the top-right plots of figures 3-5, and it is related to the trend in figure 7 that the TSVD approach systematically yields higher  $\chi_K^2$  than REGCOIL for given  $\chi_B^2$ . These aspects of the TSVD method can be understood in light of observations in [16]. In that work, it was pointed out that the SVD does not monotonically order structures in the current potential from largest to smallest scale, because  $B_{\text{normal}}$  is driven by the *gradient* of the current potential, and the gradient emphasizes small-scale structures. While the large-scale structures in the current potential generate magnetic fields that propagate efficiently from the coil to plasma surface, it is the small-scale structures in the current potential that most efficiently generate magnetic field locally. Hence, the TSVD emphasizes small-scale structures, yielding unnecessarily contorted coil shapes. This problem could possibly be avoided by using a truncated *generalized* SVD [20], using a second matrix related to  $\chi_K^2$ .

Ideally, one would like to have an automated procedure to choose an appropriate level of regularization. Such automation would be particularly useful for the ‘stage 1’ fixed-boundary plasma optimization, to identify plasma shapes that can be produced by distant coils, since the current potential may then need to be computed thousands of times. In other fields, one popular method for automatically choosing the level of regularization for an ill-posed problem is the ‘L-curve’ method [20, 25]. This method is based on the observation that for a wide class of regularized ill-posed problems, a log-log plot of the two competing criteria typically has a characteristic L shape, as can be seen in figure 7.b. In the L-curve method, one adopts the heuristic that the regularization level should be chosen to obtain a solution on the corner of this L-shaped curve. Unfortunately, the L-curve method appears to be unhelpful for the present problem of computing the current potential for stellarator coils, for two reasons. First, while an L-shaped curve does exist for W7-X (figure 7.b), the appropriate level of regularization for reasonable equilibrium reconstruction and reasonable coil curvature actually corresponds to a point on the vertical part of the L, such as the bold symbols in figure 7.b. (This will be demonstrated in the next section.) Second, for the NCSX case, the L curve in figure 7.a is quite broad, so it is hard to identify a unique corner. Instead of the L-curve heuristic, an automatic procedure for choosing the level of regularization could be to find the point on the tradeoff curve (figure 7) at which some engineering limit is reached, such as the maximum allowable current density  $K$ .






FIG. 7. (Color online) Tradeoffs in the competing criteria of  $B_{\text{normal}}$  versus current density  $K$  for the three algorithms and two geometries. As the level of regularization is scanned for each algorithm, a series of points is obtained if the regularization is discrete (NESCOIL and TSVD), while a continuous curve is obtained for REGCOIL since the regularization is continuous. A point lying above and to the right of another point is ‘dominated’ in the sense that both criteria are simultaneously improved at the latter point. Figures (a)-(b) show the tradeoff curves for the surface-average-squared quantities. Figure (c)-(d) show the tradeoff curves for the maximum quantities. Results from REGCOIL consistently dominate results from the other methods.

## VI. FREE BOUNDARY RECONSTRUCTIONS USING DISCRETE COILS

Another demonstration of the superiority of REGCOIL is shown in figure 8. Here we use the W7-X geometry from figure 2.c-d, and we compute sets of 50 discrete coils (5 per half period, as for the real experiment) using each of the three algorithms. For these calculations, the regularization parameter for each method was chosen to select points relatively close to each other in figures 7.b and d., shown in bold symbols, each suitable for a reasonable quality of plasma shape reconstruction. Then VMEC is run in free-boundary mode [26] to compute an equilibrium supported by each coilset. Figure 8.a shows that while all three coilsets do a reasonable job producing the desired plasma shape, the coilset from REGCOIL does slightly better than the others. The coilsets from NESCOIL and REGCOIL are overlaid in

figure 8.b. It can be seen that for the REGCOIL set, there is more space between the blue and magenta coils, associated with a reduction in the maximum  $K$ . In figure 8.c the coilsets from the TSVD method and REGCOIL are overlaid. It can be seen that the coils from the TSVD method have several regions with significantly higher curvature than the corresponding regions of the REGCOIL coils, particularly in the magenta and green coils. Neglecting the coil thickness, the distance of closest approach between the blue and magenta coils from REGCOIL is 0.280 m, larger (better) than the corresponding distances 0.233 m for NESCOIL and 0.174 m for TSVD. Similarly, the separation between consecutive magenta coils in adjacent half-periods is 0.331 m for REGCOIL, much improved from 0.141 m for NESCOIL and 0.199 m for TSVD. Thus, the coilsets generated from REGCOIL simultaneously do a better job of producing the desired plasma shape and maximizing space between the coils for ports and other components. This superiority is reflected in the fact that the REGCOIL solution used in figure 8 has lower  $\chi_B^2$ , lower  $\chi_K^2$ , lower  $\max(|B_{\text{normal}}|)$ , and lower  $\max(K)$  than the solutions from the other two methods.

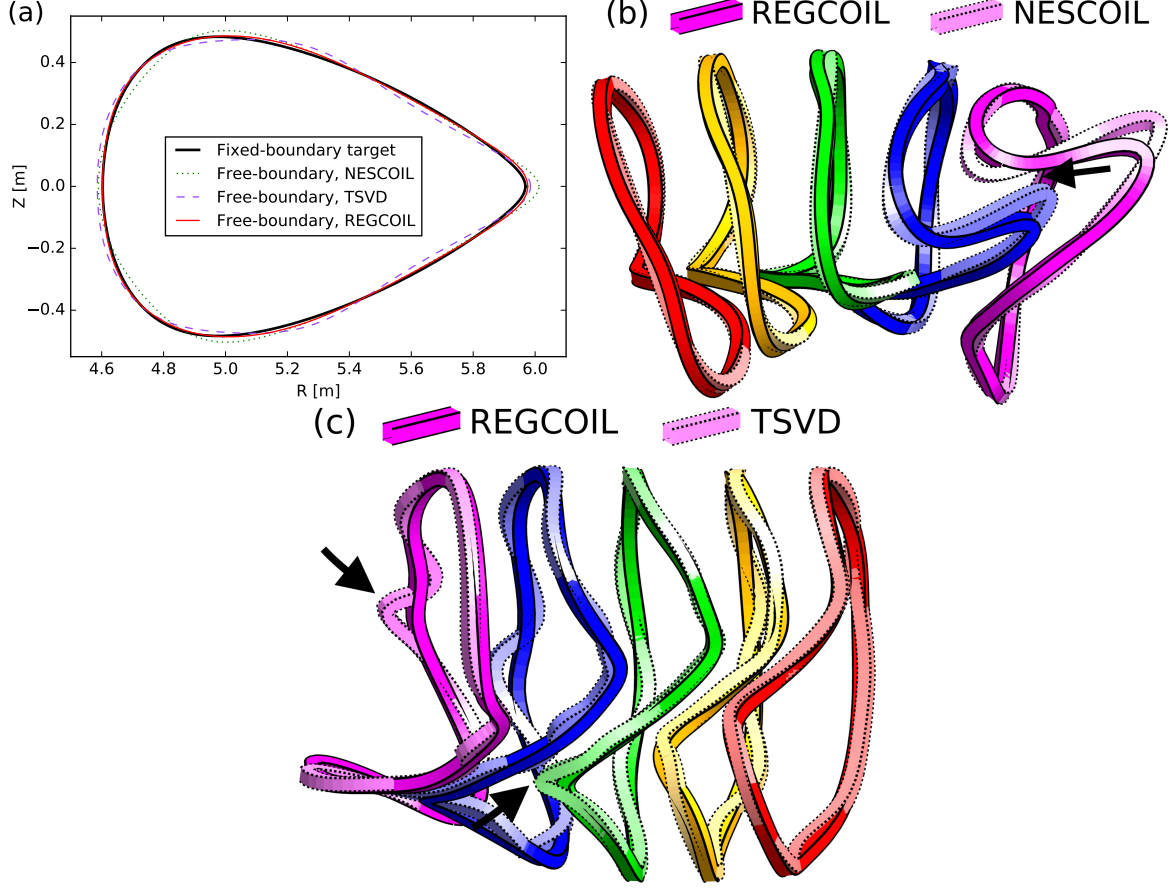


FIG. 8. (Color online) Compared to the TSVD and NESCOIL methods, REGCOIL can achieve superior reconstruction of the target plasma shape and superior coil shapes simultaneously. Here, coilsets for W7-X were computed using the three methods, with an appropriate level of regularization for a reasonable reconstruction of the target shape. The fixed-boundary target shape is shown in (a) along with free-boundary VMEC reconstructions using the three methods with sets of 50 modular coils. The REGCOIL reconstruction is closest to the target shape, associated with the lower  $\chi_B^2$  and lower  $\max(|B_{\text{normal}}|)$  in figure 7. The 5 unique coil shapes from NESCOIL and REGCOIL for this reconstruction are shown overlaid in (b). More space is available between the blue and magenta coils from REGCOIL, associated with lower  $\chi_K^2$  in figure 7. The TSVD and REGCOIL coilsets used in (a) are shown overlaid in (c), where the TSVD coils display regions of unnecessarily high curvature. The REGCOIL coils in (b) and (c) are identical, although the viewing angle is different.

Another practical advantage of REGCOIL is that for a given  $\chi_B^2$ , modular coils are less likely to bifurcate into saddle coils, indicated by contours of the current potential that do not link the plasma. This trend can be seen by comparing the middle rows of figures 3-5. In a calculation of non-planar modular coils such as this, the saddle coils are undesirable since they represent additional components to fabricate and they block port access. If one does not want saddle coils, the ability of NESCOIL and the TSVD method to reconstruct the NCSX target plasma shape is quite limited. This limitation is shown in the free-boundary reconstruction for NCSX in figure 9.a, using 3 coils per half-period. (For simplicity we take the coils to lie on equally spaced contours, carrying equal current.) For the NESCOIL and TSVD

methods, we choose the regularization parameters to give the most accurate reconstruction of the target plasma shape possible without saddle coils. In other words, we start with the maximum possible regularization, then lower the regularization until the topology of the three relevant contours is about to change. Slightly higher regularization is used for REGCOIL to smooth the coil shapes a bit. Figure 9 demonstrates that while none of the algorithms is able to achieve a perfect reconstruction of the target plasma shape, REGCOIL achieves a much better reconstruction than the other methods. The REGCOIL-derived coilset used for 9.a is shown in figure 9.b




FIG. 9. (Color online) Free-boundary NCSX reconstruction using 3 coils per half period. (a) The best possible reconstruction without saddle coils from NESCOIL and the TSVD method is not nearly as accurate as with a coilset from REGCOIL. Two half-periods of the REGCOIL coilset used for (a) are shown in (b).

## VII. COORDINATE DEPENDENCE OF THE METHODS

Since the coordinates  $\theta'$  and  $\zeta'$  parameterizing the coil surface appear explicitly in NESCOIL's regularization scheme (the Fourier series truncation), the resulting coil shapes are not the same if the coil surface is reparameterized using some other equally valid angle coordinates. This property is undesirable for several reasons. It places an extra demand on the code user, as he or she must be careful to use 'good' coordinates to parameterize the coil surface, and it is not actually clear what the best coordinates are. Since NESCOIL may well return superior coil shapes if the surface is reparameterized, it is never clear if the best solution to the optimization problem has been found. Another problem arises if NESCOIL is called within fixed-boundary plasma optimization in order to find plasma shapes consistent with good coils, as in [15]. In such an optimization, a large residual  $\chi_B^2$  or high coil complexity from NESCOIL-derived coils may not really mean that the plasma shape is hard to produce from external coils; it may instead indicate that the automatically-generated coil surface happens to be parameterized with coordinates for which NESCOIL performs suboptimally. Hence the fixed-boundary plasma optimizer may unfairly penalize such plasma shapes.

By contrast, the  $\chi_K^2$  term used for regularization in REGCOIL is defined by a coordinate-independent expression, and hence the coil shapes resulting from REGCOIL are independent of the coordinates used.

The TSVD solver implemented in the NESCOIL code [15, 17] and described here yields coordinate-dependent results, just as for standard NESCOIL. However as discussed in [16], the TSVD method can be reformulated to yield coordinate-independent results. The central idea is to replace the Fourier modes with basis functions that are defined to be orthogonal under the coordinate-independent operation  $\int d^2a(\dots)$ .

These concepts are illustrated in figure 10. We first generate a coil surface with uniform offset of 50 cm from the W7-X plasma surface using the algorithm of appendix B, (used in the BNORM code often run to prepare NESCOIL input files). This algorithm imposes a particular choice of poloidal angle on the coil surface. The coil surface can then be reparameterized using another poloidal angle, such as the angle with constant arclength  $|\partial \mathbf{r}' / \partial \theta'|$  at each  $\zeta'$ . Uniformly spaced points in these two poloidal angles are shown by the + and × symbols in figure 10.a. Figure 10.b shows the NESCOIL results for the two surface parameterizations, using Fourier modes  $|m| \leq 4$  and  $|n| \leq 4$ . The resulting 5 independent coil shapes (assuming 5 coils per half-period) for the two different surface parameterizations are plotted on top of each other to show the differences. Figure 10.c shows the TSVD solutions for the two surface parameterizations, retaining 125 singular values. Figure 10.d shows the REGCOIL solutions for the two surface parameterizations. Evidently, the coil shapes computed by NESCOIL and the TSVD approaches are substantially different when these two reasonable parameterizations of the same coil surface are used. In contrast the REGCOIL results are universal, independent of the surface parameterization.




FIG. 10. (Color online) Disturbingly, if the coil surface is reparameterized using a different poloidal angle as in (a), the coil shapes computed by (b) the original NESCOIL method and (c) the TSVD solver can change significantly. Shown here are the 5 independent W7-X coil shapes computed for a uniform 50cm offset from the plasma surface. Pale colors with dotted edges indicate the result using the original poloidal angle from the offset-surface generation algorithm; dark colors with solid edges indicate the result if the same coil surface is re-parameterized using a constant-arclength poloidal angle. (d) This pathology is not present in the REGCOIL approach, where the coils computed using the different surface parameterizations are identical (within a tiny discretization error that causes the striped pattern upon rendering.)

Parameterization-dependent behavior as in NESCOIL could also occur in nonlinear coil optimization codes, in which the coil shapes are directly optimized [5, 6, 9]. In these codes, one should ensure the dominant source of regularization is provided by a physical coordinate-independent quantity like  $\chi_K^2$  or the coil radius of curvature, not by the coordinate-dependent spline or Fourier representation of the coil shape.

## VIII. CONCLUSIONS

In this work we have compared three methods for computing the shape of stellarator coils: the classical NESCOIL procedure [4], TSVD solution of the NESCOIL least-squares problem [15, 17], and the Tikhonov regularization method we have called REGCOIL. The problem of determining the external currents that produce a given plasma shape is fundamentally ill-posed, requiring some form of regularization to obtain a meaningful and numerically well-behaved solution, and the three methods above each employ a different regularization strategy. All three methods have the advantage that they are guaranteed to provide a unique solution (for given choice of regularization parameter and coil surface parameterization), i.e there is no chance of getting ‘stuck’ in a local rather than global optimum as could occur in non-convex formulations. Also all three methods are ‘linear’ in the sense that there are no nonlinear equations involved which would require iterative solution. Hence all methods are comparably robust and fast. We find the TSVD method unnecessarily emphasizes small-scale structures in the coil shapes, for the reasons detailed in [16]. Moreover, since REGCOIL employs a more physically significant regularization approach than the other methods, we

find that **REGCOIL** has several compelling advantages:

1. It consistently yields lower surface-integrated-squared and maximum  $|B_{\text{normal}}|$  on the target plasma surface at the same time as lower surface-integrated-squared and maximum current density  $K$  on the coil winding surface (figure 7). Correspondingly, **REGCOIL** simultaneously achieves more accurate reconstruction of the desired plasma shape (figure 8.a and 9.a) while increasing the minimum distances between coils (figure 8.b-c). Due to this last fact, **REGCOIL** reduces the chances of coil-coil collisions and improves access for ports and maintenance.
2. **REGCOIL** allows finer control over the amount of regularization, since the regularization parameter is continuous rather than discrete.
3. Unlike the other methods, **REGCOIL** yields coil shapes that are independent of the arbitrary choice of coordinates used to parameterize the coil surface (figure 10), so the user need not worry if the surface has been parameterized using ‘good’ angles.

For these reasons, we contend that **REGCOIL** should be used instead of **NESCOIL** for applications in which a fast and robust method for coil calculation is needed, such as when targeting coil complexity in fixed-boundary plasma optimization, or for scoping of new stellarator geometries.

## ACKNOWLEDGMENTS

The author is grateful to Michael Drevlak for providing the W7-X plasma and coil surface data, and to Neil Pomphrey for providing the NCSX plasma and coil surface data. We also acknowledge illuminating conversations about this topic with Allen Boozer and Mike Zarnstorff, and thank Sam Lazerson for providing the **VMEC** and **NESCOIL** codes. This work was supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Science, under Award Number DE-FG02-93ER54197. Computations were performed on the Edison and Cori systems at the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

## Appendix A: Equations for **NESCOIL** and **REGCOIL**

Here we derive the equations used for the **NESCOIL** and **REGCOIL** methods. First we derive an expression for  $\chi_B^2$ , equivalent to the calculation in [4]. Next we derive an expression for  $\chi_K^2$ . The derivative of the total objective function then yields a linear system for the single-valued part of the current potential.

We first note that the surface integral of any quantity  $Q$  (such as appears in (2) and (7)) can be written

$$\int d^2a \, Q = \sum_{\ell=1}^{n_p} \int_0^{2\pi} d\theta \int_0^{2\pi/n_p} d\zeta \, NQ. \quad (\text{A1})$$

Here  $n_p$  is the number of identical toroidal periods (e.g. 5 for W7-X), the toroidal period is indexed by  $\ell$ ,  $N = |\mathbf{N}|$ , and

$$\mathbf{N} = \frac{\partial \mathbf{r}}{\partial \zeta} \times \frac{\partial \mathbf{r}}{\partial \theta} \quad (\text{A2})$$

is a non-unit-length surface normal vector.

The Biot-Savart law gives the magnetic field due to a surface current density  $\mathbf{K}$  as

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int d^2a' \frac{\mathbf{K}' \times (\mathbf{r} - \mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|^3}. \quad (\text{A3})$$

As in the main text, primed coordinates are those on the coil surface, and other primed quantities like  $\mathbf{K}'$  and  $\mathbf{r}'$  are shorthand for  $\mathbf{K}(\theta', \zeta')$  and  $\mathbf{r}(\theta', \zeta')$ . Applying (A2) and (A1), and noting

$$\begin{aligned} \mathbf{N} \times \nabla \theta &= -\partial \mathbf{r} / \partial \zeta, \\ \mathbf{N} \times \nabla \zeta &= \partial \mathbf{r} / \partial \theta, \end{aligned} \quad (\text{A4})$$



we find

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \sum_{\ell'=1}^{n_p} \int_0^{2\pi} d\theta' \int_0^{2\pi/n_p} d\zeta' \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \left( \frac{\partial \mathbf{r}'}{\partial \theta'} \frac{\partial \Phi'}{\partial \zeta'} - \frac{\partial \mathbf{r}'}{\partial \zeta'} \frac{\partial \Phi'}{\partial \theta'} \right) \times (\mathbf{r} - \mathbf{r}'). \quad (\text{A5})$$

Considering the normal component of (A5) on the plasma surface, the contribution to  $B_{\text{normal}}$  from the secular terms in (3) is

$$B_{\text{normal}}^{GI} = \frac{\mu_0}{8\pi^2} \sum_{\ell'=1}^{n_p} \int_0^{2\pi} d\theta' \int_0^{2\pi/n_p} d\zeta' \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \left( G \frac{\partial \mathbf{r}'}{\partial \theta'} - I \frac{\partial \mathbf{r}'}{\partial \zeta'} \right) \times (\mathbf{r} - \mathbf{r}') \cdot \mathbf{n} \quad (\text{A6})$$

$$\approx \frac{\mu_0}{8\pi^2} \Delta_{\theta'} \Delta_{\zeta'} \sum_{\ell'=1}^{n_p} \sum_{i'_\theta} \sum_{i'_\zeta} \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \left( G \frac{\partial \mathbf{r}'}{\partial \theta'} - I \frac{\partial \mathbf{r}'}{\partial \zeta'} \right) \times (\mathbf{r} - \mathbf{r}') \cdot \mathbf{n}. \quad (\text{A7})$$

In (A7) we have discretized the  $\theta'$  and  $\zeta'$  integrals by introducing uniform grids in these coordinates with spacing  $\Delta_{\theta'}$  and  $\Delta_{\zeta'}$ , with domains  $[0, 2\pi)$  and  $[0, 2\pi/n_p)$ , and indexed by  $i'_\theta$  and  $i'_\zeta$ . Equation (A7) corresponds to  $h/N$  in the notation of [4].

Next, we integrate (A5) by parts in  $\theta'$  and  $\zeta'$  to remove the derivatives on  $\Phi'$ , and we take the component normal to the plasma surface. Thus, we find the normal magnetic field due the single-valued current potential is, after some algebra,

$$B_{\text{normal}}^{\text{sv}}\{\Phi_{sv}\} = \frac{\mu_0}{4\pi N} \sum_{\ell'=1}^{n_p} \int_0^{2\pi} d\theta' \int_0^{2\pi/n_p} d\zeta' \Phi'_{sv} \left( \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \mathbf{N} \cdot \mathbf{N}' - \frac{3}{|\mathbf{r} - \mathbf{r}'|^5} [\mathbf{r} - \mathbf{r}'] \cdot \mathbf{N} [\mathbf{r} - \mathbf{r}'] \cdot \mathbf{N}' \right). \quad (\text{A8})$$

Considering the Fourier expansion (5), and discretizing the  $\theta'$  and  $\zeta'$  integrals as in (A7), we find the normal magnetic field driven by the single-valued current potential is

$$B_{\text{normal}}^{\text{sv}}\{\Phi_{sv}\} = \frac{1}{N} \sum_j \Phi_j g_j, \quad (\text{A9})$$

where

$$g_j(\theta, \zeta) = \Delta_{\theta'} \Delta_{\zeta'} \sum_{i'_\theta} \sum_{i'_\zeta} \left( \frac{\sin}{\cos} \right)_j (m_j \theta' - n_j \zeta') g(\theta, \zeta, \theta'_{i'_\theta}, \zeta'_{i'_\zeta}) \quad (\text{A10})$$

and

$$g(\theta, \zeta, \theta', \zeta') = \frac{\mu_0}{4\pi} \sum_{\ell'=1}^{n_p} \left( \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \mathbf{N} \cdot \mathbf{N}' - \frac{3}{|\mathbf{r} - \mathbf{r}'|^5} [\mathbf{r} - \mathbf{r}'] \cdot \mathbf{N} [\mathbf{r} - \mathbf{r}'] \cdot \mathbf{N}' \right) \quad (\text{A11})$$

are quantities introduced in [4]. Using (2) and (A1), the NESCOIL objective function can now be written

$$\chi_B^2 = n_p \Delta_\theta \Delta_\zeta \sum_{i_\theta} \sum_{i_\zeta} \left[ \sqrt{N} (B_{\text{normal}}^{\text{plasma}} + B_{\text{normal}}^{\text{external}} + B_{\text{normal}}^{GI}) + \sum_j \Phi_j \frac{g_j}{\sqrt{N}} \right]^2. \quad (\text{A12})$$

The  $\theta$  and  $\zeta$  integrals have been discretized as in (A7). Minimization of (A12) has the form of a linear least-squares problem,  $|Ax - b|^2$ , where the matrix  $A$  corresponds to  $g_j/\sqrt{N}$ , the vector  $b$  corresponds to  $-\sqrt{N}(B_{\text{normal}}^{\text{plasma}} + B_{\text{normal}}^{\text{external}} + B_{\text{normal}}^{GI})$ , and the vector  $x$  corresponds to  $\Phi_j$ . As with any such problem, it may be solved by the normal equations,  $QR$  decomposition, or SVD. The normal equations approach is typically slightly faster than the other methods but is more affected by roundoff error.

Next, we derive an expression for the current density objective  $\chi_K^2$ . Combining (1), (3), (A2), and (A4), we find the current density can be written

$$\mathbf{K}' = \frac{1}{N'} \left( \mathbf{d} - \sum_j \Phi_j \mathbf{f}_j \right), \quad (\text{A13})$$



where

$$\mathbf{d} = \frac{G}{2\pi} \frac{\partial \mathbf{r}'}{\partial \theta'} - \frac{I}{2\pi} \frac{\partial \mathbf{r}'}{\partial \zeta'} \quad (\text{A14})$$

is the contribution from the secular terms in the current potential, and

$$\mathbf{f}_j = \left[ m_j \frac{\partial \mathbf{r}'}{\partial \zeta'} + n_j \frac{\partial \mathbf{r}'}{\partial \theta'} \right] \begin{pmatrix} \cos \\ -\sin \end{pmatrix}_j (m_j \theta' - n_j \zeta') \quad (\text{A15})$$

is the contribution from the  $j$ th mode of the periodic current potential. Substituting these expressions into (7) with (A1), and discretizing the  $\theta'$  and  $\zeta'$  integrals in the same way as in (A7), we obtain

$$\chi_K^2 = n_p \Delta_{\theta'} \Delta_{\zeta'} \sum_{i'_\theta} \sum_{i'_\zeta} \frac{1}{N'} \left( \mathbf{d} - \sum_j \Phi_j \mathbf{f}_j \right)^2. \quad (\text{A16})$$

If one did not care about  $\chi_B^2$  and only wanted to minimize  $\chi_K^2$ , one would form  $\partial \chi_K^2 / \partial \Phi_j = 0$  to obtain the normal equations

$$\sum_k A_{j,k}^K \Phi_k = b_j^K, \quad (\text{A17})$$

with matrix

$$A_{j,k}^K = \Delta_{\theta'} \Delta_{\zeta'} \sum_{i'_\theta} \sum_{i'_\zeta} \frac{\mathbf{f}_j \cdot \mathbf{f}_k}{N'}, \quad (\text{A18})$$

and right-hand side

$$b_j^K = \Delta_{\theta'} \Delta_{\zeta'} \sum_{i'_\theta} \sum_{i'_\zeta} \frac{\mathbf{d} \cdot \mathbf{f}_j}{N'}. \quad (\text{A19})$$

For the combined objective function (8), setting  $0 = \partial \chi^2 / \partial \Phi_j = \partial \chi_B^2 / \partial \Phi_j + \lambda \partial \chi_K^2 / \partial \Phi_j$ , we obtain the linear system to solve for REGCOIL:

$$\sum_k A_{j,k} \Phi_k = b_j, \quad (\text{A20})$$

with matrix

$$A_{j,k} = A_{j,k}^B + \lambda A_{j,k}^K, \quad (\text{A21})$$

$$A_{j,k}^B = \Delta_\theta \Delta_\zeta \sum_{i_\theta} \sum_{i_\zeta} \frac{g_j g_k}{N}, \quad (\text{A22})$$

and right-hand side

$$b_j = b_j^B + \lambda b_j^K, \quad (\text{A23})$$

$$b_j^B = -\Delta_\theta \Delta_\zeta \sum_{i_\theta} \sum_{i_\zeta} \left( B_{\text{normal}}^{\text{plasma}} + B_{\text{normal}}^{\text{external}} + G_{\text{normal}}^{GI} \right) g_j \quad (\text{A24})$$

## Appendix B: Offset surface algorithm

Suppose we know a toroidal surface  $\mathbf{r}_p(\theta, \zeta)$ , parameterized by poloidal and toroidal angles  $(\theta, \zeta)$ , and we wish to compute a toroidal surface  $\mathbf{r}_c(\theta', \zeta')$  which is offset from  $\mathbf{r}_p$  by a uniform distance  $\Delta$ . For each pair of values of the poloidal and toroidal angles  $(\theta', \zeta')$  we can obtain  $\mathbf{r}_c(\theta', \zeta')$  as follows. Let  $\theta = \theta'$ , and solve the 1D nonlinear root-finding problem

$$\arctan\left(y(\hat{\zeta})/x(\hat{\zeta})\right) - \zeta' = 0 \quad (\text{B1})$$

for the unknown  $\hat{\zeta}$ , where  $x(\hat{\zeta})$  and  $y(\hat{\zeta})$  are the Cartesian components of the vector  $\mathbf{r}_p(\theta, \hat{\zeta}) + \mathbf{n}(\theta, \hat{\zeta})\Delta$ . Here,  $\mathbf{n}$  is the unit normal of the surface  $\mathbf{r}_p$ . Thus,  $\mathbf{r}_p(\theta', \hat{\zeta})$  is the point on the original surface such that if we move a distance  $\Delta$  in the normal direction, we arrive at the toroidal angle  $\zeta'$ . Then the desired  $\mathbf{r}_c(\theta', \zeta')$  is given by  $\mathbf{r}_p(\theta', \hat{\zeta}) + \mathbf{n}(\theta', \hat{\zeta})\Delta$ . This process can be repeated to generate as many points  $(\theta', \zeta')$  on the new surface as desired. Note that the angle  $\zeta'$  parameterizing the new surface is the standard cylindrical angle.

- 
- [1] G. Grieger, W. Lotz, P. Merkel, J. Nührenberg, J. Sapper, E. Strumberger, H. Wobig, the W7-X Team, R. Burhenn, V. Erckmann, U. Gasparino, L. Giannone, H. J. Hattfuss, R. Jaenicke, G. Kühner, H. Ringler, A. Weller, F. Wagner, and the W7-AS Team, *Phys. Fluids B* **4**, 2081 (1992).
  - [2] T. Klinger, C. Baylard, C. D. Beidler, J. Boscary, H. S. Bosch, A. Dinklage, D. Hartmann, P. Helander, H. Maassberg, A. Peacock, T. S. Pedersen, T. Rummel, F. Schauer, L. Wegener, and R. Wolf, *Fusion Eng. Des.* **88**, 461 (2013).
  - [3] F. S. B. Anderson, D. T. Anderson, A. Almagri, P. G. Matthews, J. N. Talmadge, and J. L. Shohet, *Fusion Tech.* **27**, 273 (1995).
  - [4] P. Merkel, *Nucl. Fusion* **27**, 867 (1987).
  - [5] M. Drevlak, *Fusion Tech.* **33**, 106 (1998).
  - [6] D. J. Strickler, L. A. Berry, and S. P. Hirshman, *Fusion Sci. Tech.* **41**, 107 (2001).
  - [7] D. J. Strickler, S. P. Hirshman, D. A. Spong, M. J. Cole, J. F. Lyon, B. E. Nelson, D. E. Williamson, and A. S. Ware, *Fusion Sci. Tech.* **45**, 15 (2004).
  - [8] M. C. Zarnstorff, L. A. Berry, A. Brooks, E. Fredrickson, G. Fu, S. Hirshman, S. Hudson, L. Ku, E. Lazarus, D. Mikkelsen, D. Monticello, G. H. Neilson, N. Pomphrey, A. Reiman, D. Spong, D. Strickler, A. Boozer, W. A. Cooper, R. Goldston, R. Hatcher, M. Isaev, C. Kessel, J. Lewandowski, J. F. Lyon, P. Merkel, H. Mynick, B. E. Nelson, C. Nührenberg, M. Redi, W. Reiersen, P. Rutherford, R. Sanchez, J. Schmidt, and R. B. White, *Plasma Phys. Controlled Fusion* **43** (2001).
  - [9] T. Brown, J. Breslau, D. Gates, N. Pomphrey, and A. Zolfaghari, *IEEE 26th Symposium on Fusion Engineering (SOFE)* (2015).
  - [10] L. Ku and A. H. Boozer, *Nucl. Fusion* **50**, 125005 (2010).
  - [11] L. Ku and A. H. Boozer, *Nucl. Fusion* **51**, 013004 (2011).
  - [12] M. Drevlak, F. Brochard, P. Helander, J. Kisslinger, M. Mikhailov, C. Nührenberg, J. Nührenberg, and Y. Turkin, *Contrib. Plasma Phys.* **53**, 459 (2013).
  - [13] J. Zheng, Y. Song, J. Breslau, and G. H. Neilson, *Fusion Eng. Des.* **89**, 487 (2014).
  - [14] V. Queral, *Fusion Eng. Des.* **89**, 2145 (2014).
  - [15] N. Pomphrey, L. Berry, A. Boozer, A. Brooks, R. E. Hatcher, S. P. Hirshman, L.-P. Ku, W. H. Miner, H. E. Mynick, W. Reiersen, D. J. Strickler, and P. M. Valanju, *Nucl. Fusion* **41**, 339 (2001).
  - [16] M. Landreman and A. H. Boozer, *Phys. Plasmas* **23**, 032506 (2016).
  - [17] P. Valanju, W. Miner, A. Brooks, N. Pomphrey, S. Hirshman, and L. Berry, in [http://ncsx.pppl.gov/Scientificconf/2000\\_10\\_23\\_APS/valanju.pdf](http://ncsx.pppl.gov/Scientificconf/2000_10_23_APS/valanju.pdf), *APS Meeting Abstracts*, Vol. 1 (2000) p. 1034P.
  - [18] L. Ku and A. Boozer, *Phys. Plasmas* **17**, 122503 (2010).
  - [19] A. Boozer, *J. Plasma Phys.* **81**, 515810606 (2015).
  - [20] P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion* (SIAM, 1998).
  - [21] A. S. Jacobsen, L. Stagner, M. Salewski, B. Geiger, W. W. Heidbrink, S. B. Korsholm, F. Leipold, S. K. Nielsen, J. Rasmussen, M. Stejner, H. Thomsen, M. Weiland, and the ASDEX Upgrade team, *Plasma Phys. Controlled Fusion* **58**, 045016 (2016).
  - [22] R. E. Steuer, *Multiple criteria optimization: theory, computation, and applications* (Wiley, 1986).
  - [23] A. N. Tikhonov, *Soviet Math. Dokl.* **4**, 1035 (1963).
  - [24] S. P. Hirshman and J. C. Whitson, *Phys. Fluids* **26**, 3553 (1983).
  - [25] P. C. Hansen, *SIAM Rev.* **34**, 561 (1992).
  - [26] S. P. Hirshman, W. I. van Rij, and P. Merkel, *Comp. Phys. Comm.* **43**, 143 (1986).